AI Data Hygiene Snapshot

Fictional Sample Report

PREPARED FOR

Northbridge Legal Studio (fictional)

Entry diagnostic for Agent Workflow Safety & Governance

A short operational review of current AI usage, data handling exposure, unmanaged tools, and first-week controls — completed before AI agents are granted access to real business workflows.

CLIENT
Northbridge Legal Studio (fictional)

PREPARED BY
Alpaca Data Lab

DOCUMENT TYPE
Mini Operational Snapshot

DATE OF ISSUE
06 May 2026

STATUS
Fictional sample · Demo report

DISTRIBUTION
Public demo report

IMPORTANT NOTICE

Scope and limitations of this Snapshot

This AI Data Hygiene Snapshot is a short operational review based on questionnaire-style responses and optional redacted workflow descriptions. It is designed to identify common AI usage risks and suggest practical first actions for a small team.

This Snapshot can also serve as an entry diagnostic for Agent Workflow Safety & Governance. It helps identify current AI usage patterns, unmanaged tools, data exposure, and early control gaps before a team grants AI agents access to real systems such as email, files, CRM, calendars, or document repositories.

Northbridge Legal Studio is a fictional company created for demonstration purposes. This sample report is not based on a real client, real confidential documents, real employee AI chats, real production systems, or real legal matters.

This report is not a legal opinion, cybersecurity certification, penetration test, vulnerability scan, SOC 2 assessment, ISO 27001 audit, GDPR / HIPAA / CCPA compliance opinion, breach investigation, forensic review, or substitute for qualified legal, privacy, compliance, cybersecurity, or managed IT advice.

No actual confidential client documents, employee AI chat histories, production systems, passwords, credentials, API keys, privileged materials, full customer records, or raw confidential datasets were reviewed. Any policy language or recommendation should be reviewed by appropriate internal leadership and qualified advisors before adoption.

SECTION 01

Executive Summary

Northbridge Legal Studio is a fictional 23-person boutique legal and business advisory firm serving small businesses, startup founders, consultants, creative agencies, independent software companies, and professional service firms.

Based on the sample intake, the firm currently uses or may use AI tools such as ChatGPT, Claude, Gemini, Microsoft 365 AI features, occasional AI meeting transcription tools, and browser-based AI writing assistants. Most AI usage appears to happen through personal or unmanaged accounts rather than a centralized company-managed AI platform.

The issue is not that Northbridge is using AI. The issue is that AI usage is happening faster than the firm’s data handling rules, tool visibility, and human review practices.

The next risk frontier is not only what employees paste into AI tools, but what AI agents may later be allowed to access or act on across real business systems. For Northbridge, today’s AI data hygiene gaps should be addressed before any agent is connected to inboxes, files, calendars, CRM, matter folders, or document repositories.

> *“Northbridge handles confidential and legal-adjacent client material, but employees appear to use personal AI tools without a clear approved tool list, redaction checklist, or formal review rule.”*

PROVISIONAL RISK BAND

High Priority · Score 82 / 100

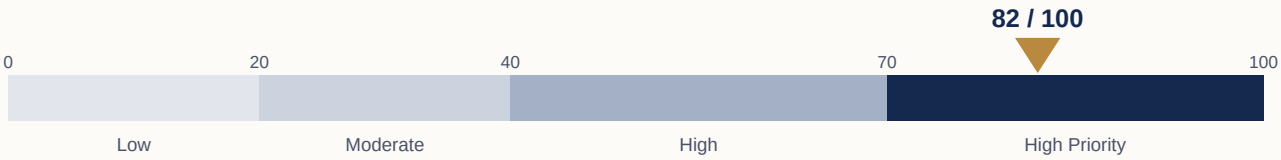

Recommended first priority. Stop or restrict unredacted client material in personal AI accounts, then create a basic AI tool list and redaction rule within seven days. Before any AI agent is connected to files, email, calendars, matter folders, CRM, or document repositories, Northbridge should define permission boundaries and human approval gates. This score is an operational prioritization signal — not a legal conclusion, certification, or evidence that a breach occurred.

Top three priorities

1. Stop personal-AI use for full unredacted client emails, contract clauses, meeting notes, or privileged / restricted materials.
2. Create a simple approved / restricted / prohibited AI tool list within seven days.
3. Publish a one-page prompt redaction checklist for client-adjacent work.

SECTION 02

AI Usage Snapshot

Organization type	Boutique legal and business advisory firm
Team size	23 people
Primary clients	Small businesses, startups, consultants, agencies, indie software, professional services
AI tools used or suspected	ChatGPT, Claude, Gemini, Microsoft 365 AI, AI meeting transcription, browser AI writing assistants
Account model	Mostly personal or mixed; no company-managed AI platform
Main AI use cases	Email rewriting, contract clause explanation, meeting note summarization, internal SOPs, marketing, proposal language
Sensitive data handled	Yes — client names, contracts, deal terms, pricing, business plans, employment data, potentially privileged information
Existing AI policy	None
Approved AI tool list	None
Prompt redaction checklist	None
Human review practice	Informal and inconsistent; no AI-specific rule
AI governance owner	None assigned
Agent workflow readiness	Early discussion only; no permission model, approval gate, or system-by-system access boundary exists yet.

SECTION 03

Why this band was assigned

Dimension	Rating	Reason
Data sensitivity	High	Client confidential documents, contracts, negotiation details, employment-related materials, and potentially privileged information.
AI tool control	High	Employees use personal accounts; no approved / restricted / prohibited AI tool list.
Employee AI behavior	High	AI used for client-adjacent tasks — email rewriting, contract explanation, meeting summaries.
Customer confidentiality	High	Client material may enter AI workflows; firm cannot yet provide a documented answer to client AI-use questions.
Human review	High	Review expected in general but no AI-specific review matrix.
Documentation maturity	High	No AI policy, redaction checklist, tool inventory, vendor review, or governance owner.
Secrets / technical risk	Moderate–Low	Not primarily software, but credentials and browser extensions remain relevant.
Access / agent risk	Moderate–High	No autonomous agents are fully deployed yet, but agentic workflows are being discussed without a clear permission model, human approval gate, or system-by-system access boundary.

Confidence: medium-high. The intake is sufficient to identify major risk themes. The score is provisional — no direct system review, tool admin review, employee AI chat review, production access review, or confidential document review was performed.

SECTION 04

Top 5 Findings

01 Personal AI accounts used for client-adjacent work

Observation. Employees use personal ChatGPT and Claude accounts for work tasks. The firm has no inventory of who uses which AI tools, whether accounts are personal or company-controlled, or how data retention and training-use settings are configured.

Why it matters. When work happens in personal AI accounts the firm has limited visibility and control. It may not know what data was entered, which tool received it, whether content was retained, or whether the account remains accessible after departure.

First control. Within seven days, create a simple AI tool inventory. Ask employees to disclose tools used and classify each as approved, restricted, or prohibited.

02 Client emails and contract clauses entered without redaction

Observation. Employees may paste full client emails or contract clauses into AI tools to rewrite tone, summarize issues, explain language, or draft responses. Redaction is inconsistent and no formal checklist exists.

Why it matters. Client emails and contracts may carry confidential names, deal terms, negotiation positions, dispute details, or privileged context. Full-text use in unmanaged AI tools creates confidentiality and client-trust concerns.

First control. Prohibit entering full unredacted client emails, contract clauses, or matter details into personal or unapproved AI accounts. Require a redacted prompt pattern for any approved AI-supported drafting.

03 Meeting notes and transcripts may create high-sensitivity exposure

Observation. AI meeting transcription or summarization tools may be used occasionally. Some meetings include client names, deal terms, legal strategy, pricing, or confidential business concerns.

Why it matters. Transcripts often carry more sensitive information than polished documents — people speak freely. Summarizing sensitive meetings through unclear tools can create unnecessary records and exposure.

First control. Create a meeting classification rule before using AI transcription. Restrict client strategy, dispute, privileged, HR, or sensitive negotiation meetings unless an approved workflow exists.

04 AI-generated client-facing work lacks a formal review rule

Observation. AI-generated email drafts and explanations may be reviewed informally, but no AI-specific rule defines when attorney, partner, manager, or subject-matter review is required before use.

Why it matters. AI-generated text can sound polished while being incomplete, misleading, or overconfident. In legal and advisory work this creates accuracy and trust risk even when no confidential data is exposed.

First control. Create a simple review matrix: light review for generic internal brainstorming; qualified attorney review for client-facing legal explanations, contract clause suggestions, dispute summaries, and high-stakes advice.

05 No AI policy, redaction checklist, or governance owner exists

Observation. Northbridge has no dedicated AI usage policy, approved AI tool list, prompt redaction checklist, AI vendor review process, AI incident path, or internal AI governance owner.

Why it matters. Without written rules, employees must guess. “Be careful” is not a workflow. Even in a small firm, unclear rules produce inconsistent behavior across attorneys, paralegals, client success, sales, and admin staff.

First control. Assign one internal owner — for example, the operations manager with partner oversight — to maintain the AI tool list, redaction checklist, stop/restrict rules, and escalation path.

> **Agent workflow implication.** These findings should be addressed before Northbridge grants any AI agent access to files, inboxes, calendars, client folders, CRM records, or document repositories. Otherwise, existing data-handling gaps may be carried into higher-impact workflows.

SECTION 05

Immediate Stop / Restrict List

These items should be stopped, paused, or restricted first. The intent is not to ban AI but to remove the highest-risk behaviors while allowing lower-risk uses to continue.

Priority	Behavior	Recommended rule
High	Full client emails pasted into personal AI accounts	Stop unless redacted and an approved tool/workflow exists
High	Contract clauses copied into AI without redaction	Restrict; require redaction and attorney review before client use
High	Privileged, dispute, HR, or financial material in external AI	Stop unless a formally approved workflow exists
High	AI-generated client-facing legal explanations sent after only light review	Restrict; require attorney review
Medium	AI meeting transcription for sensitive client or strategy meetings	Pause until a meeting classification rule exists
Medium	Unknown browser AI extensions used on sensitive documents	Restrict until reviewed
High	AI agents with access to files, email, calendars, CRM, matter folders, repositories, or document systems	Do not deploy until permission boundaries, scoped access, logging expectations, and human approval gates are defined

SECTION 06

Safe-to-Continue Uses

These uses can generally continue if no client-specific, confidential, privileged, regulated, personal, credential-related, or non-public sensitive information is included.

Use case	Safe conditions
Brainstorming blog post ideas	Public or generic topics only
Drafting generic marketing copy	No client examples unless fully anonymized
Summarizing public articles	Public source material only
Explaining public legal concepts	Treat output as educational draft, not client advice
Creating generic checklists	No client names, matter details, or confidential workflows
Rewriting non-sensitive internal announcements	No HR, compensation, or private matters
Creating generic proposal language	Remove prospect names, pricing strategy, negotiation details

> *"If a prompt would reveal a client, matter, contract, dispute, price, strategy, credential, employee issue, or non-public business fact, do not enter it into an unmanaged AI tool."*

Note. Safe-to-continue does not mean agent-ready. A workflow that is safe for manual AI drafting may still need scoped permissions, logging, and human approval before an agent can act inside business systems.

SECTION 07

Basic Data Classification

A simple AI-specific classification model. It does not need to be perfect in the first week — it only needs to help employees decide what can and cannot go into AI tools.

Class	Examples	AI-use guidance
Public	Public legal explainers, articles, website copy, templates	Generally allowed in approved tools
Internal	Internal SOPs, non-sensitive process notes, generic templates	Allowed with caution; remove unnecessary internal details
Confidential	Client and counterparty names, contract summaries, pricing, negotiation notes	Approved tools and redaction only; avoid personal AI accounts
Restricted	Privileged material, dispute strategy, sensitive employment, client financials, settlement communications	Do not use in external AI unless a formally approved workflow exists
Prohibited for AI	Passwords, API keys, private keys, credentials, full unredacted client documents, highly sensitive personal data	Never enter into AI tools

For agentic workflows. Classification should guide not only what can enter an AI tool, but also what an agent may access or act on in the system where that data lives.

SECTION 08

Prompt Redaction Checklist

Redaction rules apply to both chat-based AI tools and agent workflows where task context may include emails, documents, transcripts, CRM notes, or internal records.

Remove or replace	Example replacement
Client names	“Client A”
Client company names	“Company B”
Email addresses / phone numbers	“[redacted email]” / “[redacted phone]”
Matter and counterparty names	“Matter X” / “Vendor Y”
Contract or pricing amounts	“[redacted amount]” / “standard pricing terms”
Deal terms and negotiation positions	“commercial terms” / “a disputed position”
Dispute details	“a confidential business dispute”
Privileged context	Remove entirely unless approved workflow exists
Employee names in sensitive contexts	“Employee A”
Compensation / payroll details	Remove entirely unless approved workflow exists
Passwords, API keys, private keys, credentials	Never include
Facts unnecessary for the AI task	Remove

SECTION 09

Safe, Unsafe, and Agent-Ready Prompts

The same task can be safe or unsafe depending on what is included in the prompt and how the AI is asked to behave.

SAFER PATTERN

A client is concerned about a vendor agreement clause involving termination notice and payment timing. Rewrite the following response in a clearer, professional tone. Do not add legal advice, do not invent facts, and keep the response neutral.

[Redacted draft response here]

UNSAFE PATTERN

Here is the full client email and contract clause. Tell me what legal risks we should raise and draft the response.

Why this is unsafe: full unredacted client material is pasted into an unmanaged tool, and the AI is invited to produce legal-risk language that may be sent to a client without qualified review.

AGENT-READY PATTERN

Draft a response using the redacted summary below. Do not send the email. Return the draft for attorney review and list any assumptions separately.

Why this is agent-ready: the action boundary is explicit (draft only, no send) and the human approval gate is named.

Operating rule. For agent workflows, the operating rule should clearly state what the agent may do and when a human must approve the next action.

SECTION 10

7-Day Action List

Northbridge does not need to solve every AI governance issue in one week. The Snapshot should reduce the most obvious risk first and lay a clean foundation for any future agent workflows.

D1

Announce temporary stop / restrict rules

Send a short internal note: until rules are finalized, do not paste full client emails, contract clauses, privileged material, dispute details, credentials, or sensitive employee information into personal AI tools.

D2

Create an AI tool inventory

Ask each employee which AI tools they use, whether the account is personal / company-managed / shared / unknown, what they use it for, and whether they use it with client or internal material.

D3

Publish an approved / restricted / prohibited tool list

Approved — company-reviewed, allowed for specific use cases. **Restricted** — only for public, generic, or redacted work. **Prohibited** — not allowed for confidential, restricted, privileged, or credential-related material.

D4

Publish the redaction checklist

Distribute a one-page checklist for removing names, emails, company names, deal terms, amounts, matter names, privileged details, and credentials before AI use.

D5

Create a human-review rule

Define when attorney, partner, manager, or subject-matter review is required. Client-facing legal explanation and contract-clause suggestion: attorney review. Sensitive meeting summary: responsible attorney or manager. Generic marketing: normal content review. Non-sensitive internal notes: team lead if operationally important.

D6

Assign an AI governance owner

One person — for example, the operations manager with managing-partner oversight — should maintain the tool list, redaction checklist, stop/restrict list, employee tool inventory, and escalation path. This owner should also track which workflows are safe, restricted, or prohibited for future AI agent use.

D7

Prepare a short client-facing holding answer

“We use AI-assisted tools only as workflow support and are implementing internal rules for data handling, redaction, tool approval, and human review. We do not intentionally use unmanaged AI tools for full unredacted client-confidential materials. Substantive legal or advisory work remains subject to qualified human review.” Should be reviewed by leadership and qualified counsel before use.

SECTION 11

When to Move Beyond the Snapshot

The AI Data Hygiene Snapshot is the entry diagnostic. Some teams will use it to implement basic controls on their own. Others may need a fuller Agent Workflow Safety & Governance setup before connecting agents to real systems.

Consider Agent Workflow Governance if:

- The firm is considering agents that may access files, email, calendars, CRM, matter folders, or document repositories.
- Leadership wants a clear permission model for agent actions.
- Client-facing workflows need human approval gates.
- Clients are asking how AI is used with their data.
- The team wants written SOPs before AI adoption expands.

A governance setup may add:

- Agent-ready workflow map
- Permission model
- Human approval matrix
- Safe / restricted / prohibited workflow list
- Agent SOPs
- Governance owner responsibilities

A Full Snapshot can still be useful where deeper diagnostic detail is needed; governance setup is for teams ready to define how agents may operate inside real workflows.

QA STATUS

Human review required before client delivery. This sample demonstrates the style and structure of an entry diagnostic report. It is not evidence that any real company has adopted unsafe AI practices. The score and risk band are operational prioritization signals only.